

3c882 U.S. PTO
07/14/00

7-17-00

A


Please type a plus sign (+) inside this box → ☐


PTO/SB/05 (4/98)
Approved for use through 09/30/2000. OMB 0651-0032
Patent and Trademark Office: U.S. DEPARTMENT OF COMMERCE
Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

UTILITY PATENT APPLICATION TRANSMITTAL (Only for new nonprovisional applications under 37 C.F.R. § 1.53(b))	Attorney Docket No.	P3809
	First Inventor or Application Identifier	Enric Musoll et al.
	Title	Method and Apparatus for Improving Fetching and Dispatch of Instructions in Multithreaded Processors
	Express Mail Label No.	EL573442657US

APPLICATION ELEMENTS See MPEP chapter 600 concerning utility patent application contents.	ADDRESS TO: Assistant Commissioner for Patents Box Patent Application Washington, DC 20231
1. <input checked="" type="checkbox"/> * Fee Transmittal Form (e.g., PTO/SB/17) (Submit an original and a duplicate for fee processing)	5. <input type="checkbox"/> Microfiche Computer Program (Appendix)
2. <input checked="" type="checkbox"/> Specification [Total Pages 26] (preferred arrangement set forth below) <ul style="list-style-type: none">- Descriptive title of the Invention- Cross References to Related Applications- Statement Regarding Fed sponsored R & D- Reference to Microfiche Appendix- Background of the Invention- Brief Summary of the Invention- Brief Description of the Drawings (if filed)- Detailed Description- Claim(s)- Abstract of the Disclosure	6. Nucleotide and/or Amino Acid Sequence Submission (if applicable, all necessary) <ul style="list-style-type: none">a. <input type="checkbox"/> Computer Readable Copyb. <input type="checkbox"/> Paper Copy (identical to computer copy)c. <input type="checkbox"/> Statement verifying identity of above copies
3. <input checked="" type="checkbox"/> Drawing(s) (35 U.S.C. 113) [Total Sheets 5]	ACCOMPANYING APPLICATION PARTS
4. Oath or Declaration [Total Pages 2] <ul style="list-style-type: none">a. <input checked="" type="checkbox"/> Newly executed (original or copy)b. <input type="checkbox"/> Copy from a prior application (37 C.F.R. § 1.63(d)) (for continuation/divisional with Box 16 completed)<ul style="list-style-type: none">i. <input type="checkbox"/> DELETION OF INVENTOR(S) Signed statement attached deleting inventor(s) named in the prior application, see 37 C.F.R. §§ 1.63(d)(2) and 1.33(b).	7. <input type="checkbox"/> Assignment Papers (cover sheet & document(s))
* NOTE FOR ITEMS 9 & 13: IN ORDER TO BE ENTITLED TO PAY SMALL ENTITY FEES, A SMALL ENTITY STATEMENT IS REQUIRED (37 C.F.R. § 1.27), EXCEPT IF ONE FILED IN A PRIOR APPLICATION IS RELIED UPON (37 C.F.R. § 1.28).	8. <input type="checkbox"/> 37 C.F.R. § 3.73(b) Statement (when there is an assignee) <input checked="" type="checkbox"/> Power of Attorney
	9. <input type="checkbox"/> English Translation Document (if applicable)
	10. <input type="checkbox"/> Information Disclosure Statement (IDS)/PTO-1449 <input type="checkbox"/> Copies of IDS Citations
	11. <input type="checkbox"/> Preliminary Amendment
	12. <input checked="" type="checkbox"/> Return Receipt Postcard (MPEP 503) (Should be specifically itemized)
	13. <input checked="" type="checkbox"/> * Small Entity Statement(s) <input type="checkbox"/> Statement filed in prior application, Status still proper and desired (PTO/SB/09-12)
	14. <input type="checkbox"/> Certified Copy of Priority Document(s) (if foreign priority is claimed)
	15. <input checked="" type="checkbox"/> Other: Check for fees

16. If a **CONTINUING APPLICATION**, check appropriate box, and supply the requisite information below and in a preliminary amendment:
☐ Continuation ☐ Divisional ☒ Continuation-in-part (CIP) of prior application No: 09, 595,776
Prior application information: Examiner Not yet assigned Group / Art Unit: Not yet assigned
For **CONTINUATION** or **DIVISIONAL APPS** only: The entire disclosure of the prior application, from which an oath or declaration is supplied under Box 4b, is considered a part of the disclosure of the accompanying continuation or divisional application and is hereby incorporated by reference. The incorporation can only be relied upon when a portion has been inadvertently omitted from the submitted application parts.

17. CORRESPONDENCE ADDRESS			
<input checked="" type="checkbox"/> Customer Number or Bar Code I		<input type="checkbox"/> Correspondence address below	
			
Name		24739	
Address		PATENT TRADEMARK OFFICE	
City	State	Zip Code	
Country	Telephone	Fax	

Name (Print/Type)	Donald R. Boye	Registration No. (Attorney/Agent)	35074
Signature		Date	07/14/2000

Burden Hour Statement: This form is estimated to take 0.2 hours to complete. Time will vary depending upon the needs of the individual case. Any comments on the amount of time you are required to complete this form should be sent to the Chief Information Officer, Patent and Trademark Office, Washington, DC 20231. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. SEND TO: Assistant Commissioner for Patents, Box Patent Application, Washington, DC 20231.

**VERIFIED STATEMENT CLAIMING SMALL ENTITY STATUS
 (37 CFR 1.9(f) & 1.27(c))--SMALL BUSINESS CONCERN**

Docket Number (Optional)
P3809

Applicant or Patentee Enric Musoll et al.

Application or Patent No. NA

Filed or Issued: NA

Title: Methods and Apparatus for Improving Fetching and Dispatch of Instructions in Multithreaded Processors

I hereby declare that I am

- ☐ the owner of the small business concern identified below
☒ an official of the small business concern empowered to act on behalf of the concern identified below:

NAME OF SMALL BUSINESS CONCERN XStream Logic, Inc.

ADDRESS OF SMALL BUSINESS CONCERN 750 University Ave., Suite 270, Los Gatos, CA 95032-7698

I hereby declare that the above identified small business concern qualifies as a small business concern as defined in 13 CFR 121.12, and reproduced in 37 CFR 1.9(d), for purposes of paying reduced fees to the United States Patent and Trademark Office, in that the number of employees of the concern, including those of its affiliates, does not exceed 500 persons. For purposes of this statement, (1) the number of employees of the business concern is the average over the previous fiscal year of the concern of the persons employed on a full-time, part-time, or temporary basis during each of the pay periods of the fiscal year, and (2) concerns are affiliates of each other when either, directly or indirectly, one concern controls or has the power to control the other, or a third party or parties controls or has the power to control both.

I hereby declare that rights under contract or law have been conveyed to and remain with the small business concern identified above with regard to the invention described in:

- ☒ the specification filed herewith with title as listed above.
☐ the application identified above.
☐ the patent identified above.

If the rights held by the above identified small business concern are not exclusive, each individual, concern, or organization having rights in the invention must file separate verified statements averring to their status as small entities, and no rights to the invention are held by any person, other than the inventor, who would not qualify as an independent inventor under 37 CFR 1.9(c) if that person made the invention, or by any concern which would not qualify as a small business concern under 37 CFR 1.9(d), or a nonprofit organization under 37 CFR 1.9(e).

- Each person, concern, or organization having any rights in the invention is listed below:
☒ no such person, concern, or organization exists.
☐ each such person, concern, or organization is listed below

Separate verified statements are required from each named person, concern or organization having rights to the invention averring to their status as small entities. (37 CFR 1.27)

I acknowledge the duty to file, in this application or patent, notification of any change in status resulting in loss of entitlement to small entity status prior to paying, or at the time of paying, the earliest of the issue fee or any maintenance fee due after the date on which status as a small entity is no longer appropriate. (37 CFR 1.28(b))

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application, any patent issuing thereon, or any patent to which this verified statement is directed.

NAME OF PERSON SIGNING Mario Nemirovsky

TITLE OF PERSON IF OTHER THAN OWNER Chief Technical Officer

ADDRESS OF PERSON SIGNING 750 University Ave., Los Gatos, CA 95032-7698

SIGNATURE

DATE 7/13/00

Certificate of Express Mailing

"Express Mail" Mailing Label Number: EL573442657US

Date of Deposit: 07/14/2000

Ref: Case Docket No.: P3809

First Named Inventor: Enric Musoll et al.

Serial Number: NA

Filing Date: 07/14/2000


Title of Case: Method and Apparatus for Improving Fetching and Dispatch of Instructions in Multithreaded Processors

I hereby certify that the attached papers are being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 C.F.R. 1.10 on the date indicated above and addressed to the Commissioner of Patents and Trademarks, Washington D.C. 20231

1. Utility patent application transmittal.
2. 26 sheets of specification.
3. 5 sheets of drawings.
4. Fee transmittal.
5. Duplicate fee transmittal.
6. Declaration and Power of Attorney.
7. Verified Statement Claiming Small Entity.
8. Check for fees in the amount of \$498.00.
9. Certificate of express mailing.
10. Postcard listing contents.

Mark A. Boys

(Typed or printed name of person mailing paper or fee)


(Signature of person mailing papers or fee)

Methods and Apparatus for Improving Fetching and Dispatch of Instructions in Multithreaded Processors

by inventors Enric Musoll and Mario Nemirovsky

5

Field of the Invention

The present invention is in the area of microprocessors, and pertains more particularly to structure and function of simultaneous multithreaded processors.

10

Cross Reference to Related Documents

The present application is a continuation-in-part (CIP) of prior copending patent application S/N 09/595,776, filed on 6/16/200, which is a CIP of prior co-pending patent applications 09/216,017, filed 12/16/98, 09/240,012, filed 1/27/99, 09/273,810, filed 3/22/99 and 09/312,302 filed 5/14/99 all five of which are incorporated herein in their entirety by reference.

15

20

Background of the Invention

Multi-streaming processors capable of processing multiple threads are known in the art, and have been the subject of considerable research and development. The present invention takes notice of the prior work in this field, and builds upon that work, bringing new and non-obvious improvements in apparatus and methods to the art. The inventors have provided with this patent application an Information Disclosure Statement listing a number of published papers in the technical field of multi-

25

30

streaming processors, which together provide additional background and context for the several aspects of the present invention disclosed herein.

For purposes of definition, this specification regards a *stream* in reference to a processing system as a *hardware* capability of the processor for supporting and processing an instruction thread. A *thread* is the actual software running within a stream. For example, a multi-streaming processor implemented as a CPU for operating a desktop computer may simultaneously process threads from two or more applications, such as a word processing program and an object-oriented drawing program. As another example, a multi-streaming-capable processor may operate a machine without regular human direction, such as a router in a packet switched network. In a router, for example, there may be one or more threads for processing and forwarding data packets on the network, another for quality-of-service (QoS) negotiation with other routers and servers connected to the network and another for maintaining routing tables and the like. The maximum capability of any multi-streaming processor to process multiple concurrent *threads* remains fixed at the number of hardware *streams* the processor supports.

A multi-streaming processor operating a single thread runs as a single-stream processor with unused streams idle. For purposes of discussion, a stream is considered an *active* stream at all times the stream supports a thread, and otherwise inactive. As in various related cases listed under the cross-reference section, and in papers provided by IDS, which were included with at least one of the cross-referenced applications, superscalar processors are also known in the art. This term refers to processors that have multiples of one or more types of functional units, and an ability to issue concurrent instructions to multiple functional units. Most central processing units (CPUs) built today have more than a single functional unit of each type, and are thus superscalar processors by this

definition. Some have many such units, including, for example, multiple floating point units, integer units, logic units, load/store units and so forth. Multi-streaming superscalar processors are known in the art as well.

State-of-the-art processors typically employ pipelining, whether the processor is a single streaming processor, or a dynamic multi-streaming processor. As is known in the art, pipelining is a technique in which multiple instructions are queued in steps leading to execution, thus speeding up instruction execution. Most processors pipeline instruction execution, so instructions take several steps until they are executed. A brief description of typical stages in a RISC architecture is listed immediately below:

- a) Fetch stage: instructions are fetched from memory
- b) Decode stage: instructions are decoded
- c) Read/Dispatch stage: source operands are read from register file
- d) Execute stage: operations are executed, an address is calculated or a branch is resolved
- e) Access stage: data is accessed
- f) Write stage: the result is written in a register

Pipeline stages take a single clock cycle, so the cycle must be long enough to allow for the slowest operation. The present invention is related to the fact that there are situations in pipelining when instructions cannot be executed. Such events are called hazards in the art. Commonly, there are three types of hazards:

- a) Structural
- b) Data
- c) Control

A structural hazard means that there are not adequate resources (e.g., functional units) to support the combination of instructions to be executed in the same clock cycle. A data hazard arises when an instruction depends on the result of one or more previous instructions not resolved. Forwarding or bypassing techniques are commonly used to reduce the impact of data

hazards. A control hazard arises from the pipelining of branches and other instructions that change the program counter (PC). In this case the pipeline may be stalled until the branch is resolved.

Stalling on branches has a dramatic impact on processor
5 performance (measured in instructions executed per cycle or IPC). The longer the pipelines and the wider the superscalar, the more substantial is the negative impact. Since the cost of stalls is quite high, it is common in the art to predict the outcome of branches. Branch predictors predict branches as either *taken* or *untaken* and the target address. Branch
10 predictors may be either static or dynamic. Dynamic branch predictors may change prediction for a given branch during program execution.

A typical approach to branch prediction is to keep a history for each branch, and then to use the past to predict the future. For example, if a
15 given branch has always been taken in the past, there is a high probability that the same branch will be taken again in the future. On the other hand, if the branch was taken 2 times, not taken 5 times, taken again once, and so forth, the prediction made will have a low confidence level. When the prediction is wrong, the pipeline must be flushed, and the pipeline control must ensure that the instructions following the wrongly guessed branch are
20 discarded, and must restart the pipeline from the proper target address. This is a costly operation.

Multistreaming processor architectures may be either fine-grained or coarse-grained. Coarse-grained multistreaming processors typically have multiple contexts, which are used to cover long latencies arising, for
25 example, due to cache misses. Only a single thread is executing at a given time. In contrast, fine-grained multistreaming technologies such as Dynamic Multi-Streaming (DMS), which is a development of XStream Logic, Inc., with which the present inventors are associated, allow true

multi-tasking or multistreaming in a single processor, concurrently executing instructions from multiple distinct threads or tasks. DMS processors implement multiple sets of CPU registers or hardware contexts to support this style of execution.

5 Increasing the relative amount of instruction level parallelism (ILP) for a processor reduces data and control hazards, so applications can exploit increasing number of functional units during peak levels of parallelism, and Dynamic Multi-Streaming (DMS) hardware and techniques within today's general-purpose superscalar processors significantly improves performance
10 by increasing the amount of ILP, and more evenly distributing it within workload. There are still occasions, however, for degraded performance due to poor selection in fetching and dispatching instructions in a DMS processor.

 What is clearly needed is improved methods and apparatus for
15 utilizing hit/miss prediction in pipelines in dynamic multi-streaming processors, particularly at the point of fetch and dispatch operations.

Summary of the Invention

20 In a preferred embodiment of the present invention, in a multi-streaming processor, a system for fetching instructions from individual ones of the multiple streams to a pipeline is provided, comprising a fetch algorithm for selecting from which stream to fetch instructions, and a
25 branch predictor for forecasting whether a branch alternative of a branch instruction will be taken. The prediction by the branch predictor is used by the fetch algorithm in determining from which stream to fetch.

In some embodiments a prediction that a branch will not be taken precipitates no change in the fetching process. Also, a prediction that a branch will be taken results in switching fetching to a different stream.

5 In some cases the branch predictor determines a probability that a branch alternative will be taken, and the probability is used by the fetch algorithm in determining from where to fetch next instructions. In other embodiments the forecast of the branch predictor is also used by a dispatch algorithm in selecting instructions from the pipeline to dispatch to functional units.

10 In another aspect of the invention, in a multi-streaming processor, a system for fetching instructions from individual ones of the multiple streams to a pipeline is provided, comprising a fetch algorithm for selecting from which stream to fetch instructions, and one or both of a branch predictor for forecasting whether a branch alternative of a branch instruction
15 will be taken, or a hit-miss predictor for forecasting whether instructions will hit or miss a data cache. In this embodiment the prediction by either or both of the predictors is used by the fetch algorithm in determining from which stream to fetch.

20 In some embodiments a prediction that a branch will not be taken or that an instruction will hit the data cache precipitates no change in the fetching process. Also in some embodiments a prediction that a branch will be taken or that an instruction will miss a data cache results in switching fetching to a different stream.

25 In some cases one or both of the branch predictors determine a probability that a branch alternative will be taken or that an instruction will miss the cache, and the probability is used by the fetch algorithm in determining from where to fetch next instructions. Also, the forecast of one or both predictors may be also used by a dispatch algorithm in selecting instructions from the pipeline to dispatch to functional units.

In yet another aspect of the invention a multi-streaming processor is provided, comprising a fetch algorithm for selecting from which stream to fetch instructions, and a branch predictor for predicting whether jumps proposed by branch instructions will be taken or not. A prediction by the branch predictor is used by the fetch algorithm in determining from where stream to fetch.

In some of these embodiments a prediction that a branch will not be taken precipitates no change in the fetching process, and a prediction that a branch will be taken results in switching fetching to a different stream. The branch predictor may determine a probability for whether a branch will be taken, and the probability is used by the fetch algorithm in determining from where to fetch next instructions. In some cases the forecast of the branch predictor is also used by a dispatch algorithm in selecting instructions from the pipeline to dispatch to functional units.

In still another embodiment a multistreaming processor is provided, comprising multiple physical streams for running individual threads, a data cache, a fetch algorithm for selecting from which stream to fetch instructions, and one or both of a branch predictor for forecasting whether a branch alternative of a branch instructions will be taken, or a hit-miss predictor for forecasting whether instructions will hit or miss a data cache. The prediction by either or both of the predictors is used by the fetch algorithm in determining from which stream to fetch. In some embodiments a prediction that a branch will not be taken or that an instruction will hit the data cache precipitates no change in the fetching process, while in others a prediction that a branch will be taken or that an instruction will miss a data cache results in switching fetching to a different stream.

In some cases one or both of the branch predictors determine a probability that a branch alternative will be taken or that an instruction will

miss the cache, and the probability is used by the fetch algorithm in determining from where to fetch next instructions, and the forecast of one or both predictors may be used by a dispatch algorithm in selecting instructions from the pipeline to dispatch to functional units.

5 Methods for practicing the invention are taught as well, and, in the various embodiments described in enabling detail below, for the first time apparatus and methods are applied to multistreaming processors to significantly improve their performance.

Brief Description of the Drawing Figures

Fig. 1a is a simplified diagram of a pipeline in an embodiment of the present invention.

15 Fig. 1b shows the pipeline of Fig. 1a after a cycle.

Fig. 1c shows the pipeline of Fig. 1a and 1b after another cycle.

Fig. 1d shows the pipeline of Fig. 1a, 1b and 1c after yet another cycle.

20 Fig. 2 is a schematic diagram associating predictors with streams in an embodiment of the present invention.

Fig. 3 is a schematic showing predictors for different levels in cache.

Fig. 4 is a schematic illustrating benefits of the technique in embodiments of the invention.

Fig. 5 is a depiction of a program counter sequence.

Description of the Preferred Embodiments

Fig. 1a is a simplified diagram of a pipeline in a dynamic, multi-streaming (DMS) processor according to an embodiment of the present invention. In this simplified view the pipeline has seven stages, which are fetch, decode, read, dispatch, execute, access and write. These are the same as described in the background section above, except for the separation of read and dispatch in Fig. 1a to illustrate the functions. Dispatch is important in the present invention in that the present invention adds intelligence to Dispatch, improving the performance of the processor. The fetch stage in the pipeline fetches instructions into the pipeline from the multiple streams, and in an embodiment of the present invention is capable of selective fetching.

Although there is no requirement in operating processors that there be instructions at each stage of a pipeline, it is often true that this is the case, and the inventors choose to illustrate each stage as occupied by a single instruction to avoid confusion in description. In many cases there will a plurality of instructions at various stages, or none at all.

In Fig. 1a the instructions in the pipeline are arbitrarily indicated as instructions A through G, at successive stages in the pipeline at one point in time. Fig. 1b shows the pipeline of Fig. 1a one cycle later. Note that instruction A has moved from fetch to decode, and the other instructions shown in Fig. 1a have moved one stage forward as well. Also, a new instruction, H, has entered the pipeline at the fetch stage.

Fig. 1c shows the same pipeline one cycle later. All instructions have moved forward one further stage, and a new instruction I has entered the pipeline at the fetch stage. Fig. 1d shows the same pipeline after yet another cycle, at which point in time the instructions have moved forward yet again, and yet another instruction J has entered the pipeline.

Note that after the fourth cycle, instruction A has moved from fetch to dispatch. Assume for the sake of this example that instruction A is a load

instruction for loading a data value from cache. If this is the case, there will be some probability as to whether the particular data is in cache or not. In the art this is known as the hit/miss probability. If the data is in the cache, the system scores a hit. If not, the system scores a miss.

5 The combination of hit/miss probability for load operations with pipelined architecture has significance for processor efficiency, because, in the conventional case the general sequence of instructions in the pipeline will be from a single thread, and will typically be related in that many instructions following a load instruction may depend upon the result of
10 whatever instruction is to use the data loaded. That is, until the resolution of whatever instruction is to use the data loaded, many following instructions cannot be executed, except in some cases, on a speculative basis.

 Conventional processors simply assume a hit when a load
15 instruction enters a pipeline. If the load is a *miss*, however, once the load instruction is executed, then it may take a number of cycles for the needed data, not in cache, to be loaded from memory. And, unfortunately, the miss will not be apparent until the load instruction is dispatched and executed. The following instructions have to stall until the data is loaded and the
20 instruction(s) depending on the data are executed.

 The present inventors provide apparatus and methods for reducing the impact of data cache misses in multithreaded architectures. The technique consists of predicting, for each of the threads running in the multiple streams of the DMS, whether the next access to the data cache will
25 result in a miss. If this is the case, then (generally):

- The stream can be given a lower priority when deciding, in the fetch stage, from which stream to fetch, and

- The dependent instructions of the instruction that accesses the data cache can be more efficiently dispatched to the functional units (FU's) in the dispatch stage.

5 This new apparatus and technique improves the performance of a multistreaming processor in the fetching and dispatching of instructions.

Fetching With Hit-Miss Prediction

10 The new technique takes advantage of the fact that, in DMS processor, as instructions are fetched to the pipeline from individual ones of the streams, there is freedom in choosing a fetching policy or algorithm that will select, on a cycle-by-cycle basis, from which stream instructions are to be fetched.

15 In a multistreaming architecture, without the technique proposed here, a typical event that causes a thread switch is a data cache miss. Since the required data may take several cycles to be available (the exact number depending on where the data really resides in the memory hierarchy of the processor), the thread that missed the data cache may be switched out since
20 the dependent instructions of the instruction that missed most likely will not execute due to the dependencies on the data. Thus, more work can be done by fetching and executing instructions from another thread. In this case, the instructions following the one that missed, and that have already been
25 fetched, will need to be flushed out, thus degrading the performance of the processor with respect to the case in which useful instructions had been fetched.

 If the fact that an instruction will miss the data cache could be known early in the process the fetching of instructions that might eventually be flushed may be avoided by fetching, instead of the instructions following

the instruction that missed the data cache, instructions from another stream, improving the likelihood that the fetched instructions may be quickly executed. Thus, a fetching algorithm, in an embodiment of the present invention, may take into account, for all the streams, the predictions on whether the next access will miss the data cache, and fetch from the stream running a thread that is most likely to have its instructions executed and committed.

There already exist in the art a variety of implementations for hit-miss predictors. The goal, however, is always the same: to predict with the highest accuracy both the hits and misses to the data cache. Moreover, a desirable property of such a predictor is to be able to predict the next access to the data cache as soon as possible so that fewer instructions (that would eventually be flushed out) will enter the pipeline.

The technique taught herein can be improved by associating a confidence level to the prediction. The predictor, in one embodiment of the invention, operating at the fetch stage, in addition to predicting also generates this confidence level value. The confidence level helps the fetching algorithm, for example, in cases in which two or more predictors predicted a miss in the data cache and one is selected to be switched out. In this case, the stream with higher confidence level will be selected.

Fig. 2 is a schematic diagram of a fetching algorithm in a multistreaming architecture. The algorithm decides from which stream(s) to fetch based on cache hit/miss predictors associated to each of the streams. In Fig. 2 a predictor is associated with streams 1, 2, and so on through stream S. Thus, theoretically, instructions from up to S streams (S being the maximum number of streams supported by the multistreaming architecture) can be simultaneously fetched every cycle. In reality, however, the fetching algorithm might be restricted to fetch instructions from P streams ($P < S$) due to implementation restrictions (for example, availability of instruction

cache ports). Moreover, the fetching algorithm might select from which streams to fetch based on other information (for example, confidence on the branch prediction for each stream, thread priorities, state of the pipeline, etc.)

5 So far, we have mentioned predictors of hit/miss for the data cache. Note that the data cache might be implemented for performance reasons in different levels (the first level - L1- being the closest to the processor core). In alternative embodiments of the invention different hit/miss predictors may exist for each of the data cache levels.

10 The fetching algorithm in alternative embodiments of the present invention may base selection of instructions to be fetched on the prediction for the second level - L2 - of data cache since, in most processor systems, a miss in the second level of cache is very costly in number of cycles (whereas the penalty of a miss in the L1 is comparatively relatively small).

15

Fetching Discrimination by Branch Prediction

20 As was described in some detail above in the "Background" section, a control hazard arises from the pipelining of branches and other instructions that change the program counter (PC). In this case the pipeline may be stalled until the branch is resolved. The description above relates in particular to the probability of whether instructions in the pipeline will hit or miss the data cache; that is, whether the data needed to execute these
25 instructions may or may not be in the cache. In the present case discrimination is accomplished by branch prediction, rather than cache hit-miss prediction.

 Stalling on branches has a dramatic impact on processor performance (measured in instructions executed per cycle or IPC). The

longer the pipelines and the wider the superscalar in a processor, the more substantial is the negative impact. Since the cost of stalls is quite high, it is common in the art in regard to single-streaming processors to predict the outcome of branches. Branch predictors predict whether a branch instruction will be taken, and may also indicate a confidence level for branch instructions and the target address if the branch is taken. Branch predictors may be either static or dynamic. Dynamic branch predictors may change prediction for a given branch during program execution.

A typical approach to branch prediction is to keep a history for each branch, and then to use the past to predict the future. For example, if a given branch has always been taken in the past, there is a high probability that the same branch will be taken again in the future. On the other hand, if the branch was taken 2 times, not taken 5 times, taken again once, and so forth, the prediction made will have a low confidence level. When the prediction is wrong, the pipeline must be flushed, and the pipeline control must ensure that the instructions following the wrongly guessed branch are discarded, and must restart the pipeline from the proper target address. This is a costly operation.

To further illustrate, Fig. 5 is a generic diagram of a program counter (PC) sequence for a specific thread, showing instructions 0 through 9 in sequence. Instruction 3 is a Branch instruction, specifically that if x is less than 2, jump to instruction 9, and if not, continue with the thread sequence at instruction 4. In a pipelined processor, when Br instruction 3 is fetched, since it will be some several cycles at least before it is dispatched to functional units and resolved, it would be good to know the likelihood as to whether the branch will be taken. If, at the time of fetching the branch instruction into the pipeline, a branch predictor is employed, and the likelihood that the branch will be taken is found to be high, and the target address is 9, a decision can be made to begin to fetch new instructions into

the pipeline at instruction 9. If the likelihood is low, then new instructions may be fetched into the pipeline sequentially, and processor performance may be considerably improved by use of the branch predictor.

The inventors have provided, in a preferred embodiment of the present invention comprising a multi-streaming processor, a system in which a branch predictor is associated with each stream of the processor to predict, to the greatest possible degree, whether a branch will be taken, and in a preferred embodiment, the confidence level of the prediction. Output from the branch predictors is fed as input to a fetching algorithm to aid in determining from which stream to fetch instructions into the pipeline.

Fig. 2 described above in the case of hit-miss prediction may also serve to illustrate the instant case for branch prediction. Again S streams are indicated, and a predictor is associated with each stream. The predictor in this case is a branch predictor, rather than the hit-miss predictor described above. As branch instructions are fetched and enter the pipeline in the multi-streaming processor, the branch predictor associated with each stream determines the probability that the branch will enter the pipeline. The predictions are fed as input to the fetching algorithm as shown, and the fetching algorithm may be structured to use this input, and perhaps other input as well, in making important decisions. In this case, a low probability that a branch will be taken allows the processor to continue with whatever fetching intelligence is currently in use. A high probability that a branch may be taken, if no target address is predicted, may be used to cause the fetching algorithm to begin fetching from a different stream than the stream from which the branch instruction was taken. If the probability that a branch will be taken is high, and a target address is predicted for the branch, further instructions may be fetched beginning from the target address.

For a given branch, a branch predictor predicts that a branch will be taken or not taken, and also may generate a confidence level of the

prediction. In a preferred embodiment the confidence level (probability) is given by a number p between 0 (about half of the time is true) to 1 (certainty). A value close to unity means it is highly likely that the prediction will become true. In a preferred embodiment a confidence-level field (CLF) of N bits is added to the branch predictor.. The N bits are a digitalization of p . For example, if $N = 1$, $CLF = 0$ if the confidence level is low and one otherwise; for $N = 2$ there are 4 levels of confidence, say, from certainty to the lowest level. The fetching algorithm makes a decision based on the value of CLF such as to fetch branch instructions from streams with the highest CLF. When a branch with low value of CLF is resolved, if no fetching from that stream has taken place following the offending branch, the CLF for that branch could be upgraded to a higher value. Meanwhile, instructions from other streams were fetched maintaining resources occupied, and avoiding the risk of stalling the pipeline.

Dispatch With Hit-Miss Prediction

The technique of having a data cache hit/miss predictor is also useful in the process of deciding, at the dispatch stage in the pipeline, which instructions are to be extracted from the instruction queue (if any) and sent to the functional units (FUs) for execution.

In current art, when an instruction (henceforth called a producer) generates a read access to the data cache, the latency of the result is not known until the data cache is accessed and the hit/miss outcome is determined. The dispatch of a dependent instruction (henceforth termed a consumer) on the data generated by the producer can follow two policies:

- a) Dispatch the instruction only when it is guaranteed that the data will be available.
- b) Dispatch the instruction assuming that the producer will hit in the first level of the data cache.

5

Policy (b), then, dispatches the consumer instruction speculatively (a hit is always assumed for the producer instruction since the hit ratio in a cache is usually very high). If the consumer instruction arrives to the FU and the data is still not available, the instruction has to either stall at the FU
10 or be rescheduled for dispatch in a later cycle (this option will allow other non-dependent instructions to be dispatched to the FU). In any case, both options degrade the performance of the processor.

Policy (a) provides the lowest performance since the consumer instruction might be unnecessarily stalled before it is dispatched. The
15 producer instruction will be dispatched as soon as the producer hits in the data cache or, in case it misses, when the missing data arrives from the next level of memory hierarchy. On the other hand, this policy provides the simplest implementation, since no re-scheduling will occur.

In an embodiment of the present invention a hit/miss predictor
20 enhances the performance of policy (b) by predicting whether the producer will hit in the data cache. Thus, the consumer instructions of a producer that is predicted to miss in the data cache will be dispatched following policy (a). If the producer instruction is predicted to hit, then the dispatch policy is (b). In this case, however, the re-scheduling logic is still needed in
25 case the prediction is incorrect. Only in the case in which the prediction is a hit but the real outcome is a miss, the consumer instructions will need to be either stalled at the FUs or re-scheduled.

In general, the hit/miss predictor operating at the dispatch level optimizes the dispatch of consumer instructions by predicting the latency of

the data. If a hit in the L1 is predicted, the latency of the data is predicted to be the latency of the L1 cache. If a miss is predicted, the predicted latency of the data depends on whether more levels of cache exist and on whether a hit/miss predictor exists for each of these levels. If, for example, two levels of cache exist and the hit/miss outcome of the L2 is also predicted, the predicted latency of the data is computed as shown in Fig. 3 (Note: the necessary cycles, if any, to bring the data from the output of the cache to the input of the functional unit where the consumer will be executed need to be added to the predicted latency of the data).

The benefits of a hit/miss predictor for dispatch logic are not restricted to multistreaming processors only, but in a multistreaming processor where the technique has larger benefits than in a conventional (single-streaming) processor architecture. In a conventional processor having a data hit/miss predictor, when a data cache miss is predicted, no instructions (in case of an in-order dispatch engine), or only those that do not depend on the missing data (in case of an out-of-order dispatch engine) can execute. In any case, the processor resources might be idle for several cycles until the missing data is available. In multistreaming processors those idle cycles can be used to execute other instructions from other threads since they do not depend on the missing data. Thus, for a multistreaming processor, the benefits of a data cache hit/miss predictor are twofold as shown in Figure 3.

Discrimination at Dispatch by Branch Prediction

Discrimination at the dispatch stage in a multi-streaming processor using hit-miss prediction is described above. Branch prediction can be used at the dispatch stage as well to improve processor performance. In a

preferred embodiment, wherein branch prediction is used at the fetch stage as input to a fetch algorithm as described above, for every branch that enters the pipeline a there will be a prediction, possibly with an attached probability, for the branch instruction. This information may be retained and passed from the fetch algorithm to a dispatch algorithm, and used in selective dispatching of instructions fetched right after the branch instruction. In one simple case, for example, the instructions following a high probability branch instructions may be given preference in dispatch versus other instructions.

In an alternative embodiment, wherein fetch discrimination is not employed, discrimination at the dispatch stage may still be used. It will be apparent to the skilled artisan, once given the teachings herein, that hit-miss and branch prediction may be done singly or in tandem at either or both of fetch and dispatch stages in a pipelined processor.

In alternative embodiments of the invention the prediction can be done differently at the fetch and dispatch stages (i.e. using different information on which to base the prediction and/or using a different prediction algorithm). As an example, the hit-miss prediction at the dispatch stage could use the program counter (PC) address of the consumer instruction (since the instruction has already been decoded and its PC is known) and could follow an algorithm similar to the prediction schemes used in branch prediction. The prediction at the fetch stage may use another type of address (cache line, for example) or other non-address information.

The prediction algorithm in different embodiments may vary depending on the workload that the processor has to efficiently support. For traditional applications, like Windows programs or SPEC benchmarks, similar algorithms to those used in branch prediction may produce the desired prediction accuracy in both hits and misses for the hit-miss case. For other types of workloads, like packet processing applications in network

processors, the predictors can take advantage of additional information, like the flow number to which the packet being processed belongs (the data cache accesses performed by the processing of the first packet(s) of a new flow most likely will miss).

5 It will be apparent to the skilled artisan that there are many alterations that might be made in the embodiments of the invention taught herein without departing from the spirit and scope of the invention. The predictors may be implemented in various ways, for example, and different actions may be taken based on assigned probabilities. Further, the
10 predictors may be used at different levels in a pipeline. For example, a predictor may have input from a decode stage, and output to a fetch algorithm. Further, the mechanisms to accomplish different embodiments of the invention may be implemented typically in either hardware or software. There are similarly many other alterations that may be made
15 within the spirit and scope of the invention. The invention should be accorded the scope of the claims below.

What is claimed is:

1. In a multi-streaming processor, a system for fetching instructions from individual ones of the multiple streams to a pipeline, comprising:
 - 5 a fetch algorithm for selecting from which stream to fetch instructions; and
 - a branch predictor for forecasting whether a branch alternative of a branch instructions will be taken;
 - wherein the prediction by the branch predictor is used by the fetch
 - 10 algorithm in determining from which stream to fetch.
2. The system of claim 1 wherein a prediction that a branch will not be taken precipitates no change in the fetching process.
- 15 3. The system of claim 1 wherein a prediction that a branch will be taken results in switching fetching to a different stream if no target address is provided by the predictor.
4. The system of claim 1 wherein the branch predictor determines a
20 probability that a branch alternative will be taken, and the probability is used by the fetch algorithm in determining from where to fetch next instructions.
5. The system of claim 1 wherein the forecast of the branch predictor is also
25 used by a dispatch algorithm in selecting instructions from the pipeline to dispatch to functional units.
6. In a multi-streaming processor, a system for fetching instructions from individual ones of the multiple streams to a pipeline, comprising:

a fetch algorithm for selecting from which stream to fetch instructions; and

one or both of a branch predictor for forecasting whether a branch alternative of a branch instructions will be taken, or a hit-miss predictor for forecasting whether instructions will hit or miss a data cache;

wherein the prediction by either or both of the predictors is used by the fetch algorithm in determining from which stream to fetch.

7. The system of claim 6 wherein a prediction that a branch will not be taken or that an instruction will hit the data cache precipitates no change in the fetching process.

8. The system of claim 6 wherein a prediction that a branch will be taken or that an instruction will miss a data cache results in switching fetching to a different stream if no target address is provided by the predictor.

9. The system of claim 6 wherein one or both of the branch predictors determine a probability that a branch alternative will be taken or that an instruction will miss the cache, and the probability is used by the fetch algorithm in determining from where to fetch next instructions.

10. The system of claim 6 wherein the forecast of one or both predictors is also used by a dispatch algorithm in selecting instructions from the pipeline to dispatch to functional units.

11. A multi-streaming processor comprising:

a fetch algorithm for selecting from which stream to fetch instructions; and

a branch predictor for predicting whether jumps proposed by branch instructions will be taken or not;

wherein a prediction by the branch predictor is used by the fetch algorithm in determining from which stream to fetch.

5

12. The processor of claim 11 wherein a prediction that a branch will not be taken precipitates no change in the fetching process.

10

13. The processor of claim 11 wherein a prediction that a branch will be taken results in switching fetching to a different stream if no target address is provided by the predictor.

15

14. The processor of claim 11 wherein the branch predictor determines a probability for whether a branch will be taken, and the probability is used by the fetch algorithm in determining from where to fetch next instructions.

20

15. The processor of claim 11 wherein the forecast of the branch predictor is also used by a dispatch algorithm in selecting instructions from the pipeline to dispatch to functional units.

16. A multistreaming processor, comprising:

multiple physical streams for running individual threads;

a data cache;

a fetch algorithm for selecting from which stream to fetch

25

instructions; and

one or both of a branch predictor for forecasting whether a branch alternative of a branch instructions will be taken, or a hit-miss predictor for forecasting whether instructions will hit or miss a data cache;

wherein the prediction by either or both of the predictors is used by the fetch algorithm in determining from which stream to fetch.

5 17. The processor of claim 16 wherein a prediction that a branch will not be taken or that an instruction will hit the data cache precipitates no change in the fetching process.

10 18. The processor of claim 16 wherein a prediction that a branch will be taken or that an instruction will miss a data cache results in switching fetching to a different stream if no target address is provided by the predictor.

15 19. The processor of claim 16 wherein one or both of the branch predictors determine a probability that a branch alternative will be taken or that an instruction will miss the cache, and the probability is used by the fetch algorithm in determining from where to fetch next instructions.

20 20. The processor of claim 16 wherein the forecast of one or both predictors is also used by a dispatch algorithm in selecting instructions from the pipeline to dispatch to functional units.

21. In a multi-streaming processor, a method for fetching instructions from individual ones of multiple streams as instruction sources to a pipeline, comprising the steps of:

25 (a) on loading a branch instruction, making a prediction by a branch predictor as to whether a branch will be taken or not; and

(b) if the prediction is that the branch will be taken, altering the source of the fetch if no target address is provided by the predictor.

22. The method of claim 21 wherein the predictor determines a probability, and the probability is used in determining fetch source.

23. In a multi-streaming processor having a data cache, a method for
5 fetching instructions from individual ones of multiple streams as instruction sources to a pipeline, comprising the steps of:

(a) on loading an instruction, making a prediction by one or both of
a branch predictor as to whether a branch will be taken if the instruction is a
branch instruction, or by a hit-miss predictor as to whether the instruction
10 will hit the data cache; and

(b) discriminating from which stream to continue to fetch according
to prediction made.

24. The method of claim 23 wherein the predictor or predictors determine a
15 probability, and the probability is used in determining fetch source.

Abstract of the Disclosure

In a multi-streaming processor, a system for fetching instructions from individual ones of multiple streams to an instruction pipeline is provided, comprising a fetch algorithm for selecting from which stream to fetch an instruction, and one or more predictors for forecasting whether a load instruction will hit or miss the cache or a branch will be taken. The prediction or predictions are used by the fetch algorithm in determining from which stream to fetch. In some cases probabilities are determined and also used in decisions, and predictors may be used at either or both of fetch and dispatch stages.

Fetch	INST A
Decode	INST B
Read	INST C
Dispatch	INST D
Execute	INST E
Access	INST F
Write	INST G

Fig. 1a

Fetch	INST H
Decode	INST A
Read	INST B
Dispatch	INST C
Execute	INST D
Access	INST E
Write	INST F

Fig. 1b

Fetch	INST I
Decode	INST H
Read	INST A
Dispatch	INST B
Execute	INST C
Access	INST D
Write	INST E

Fig. 1c

Fetch	INST J
Decode	INST I
Read	INST H
Dispatch	INST A
Execute	INST B
Access	INST C
Write	INST D

Fig. 1d

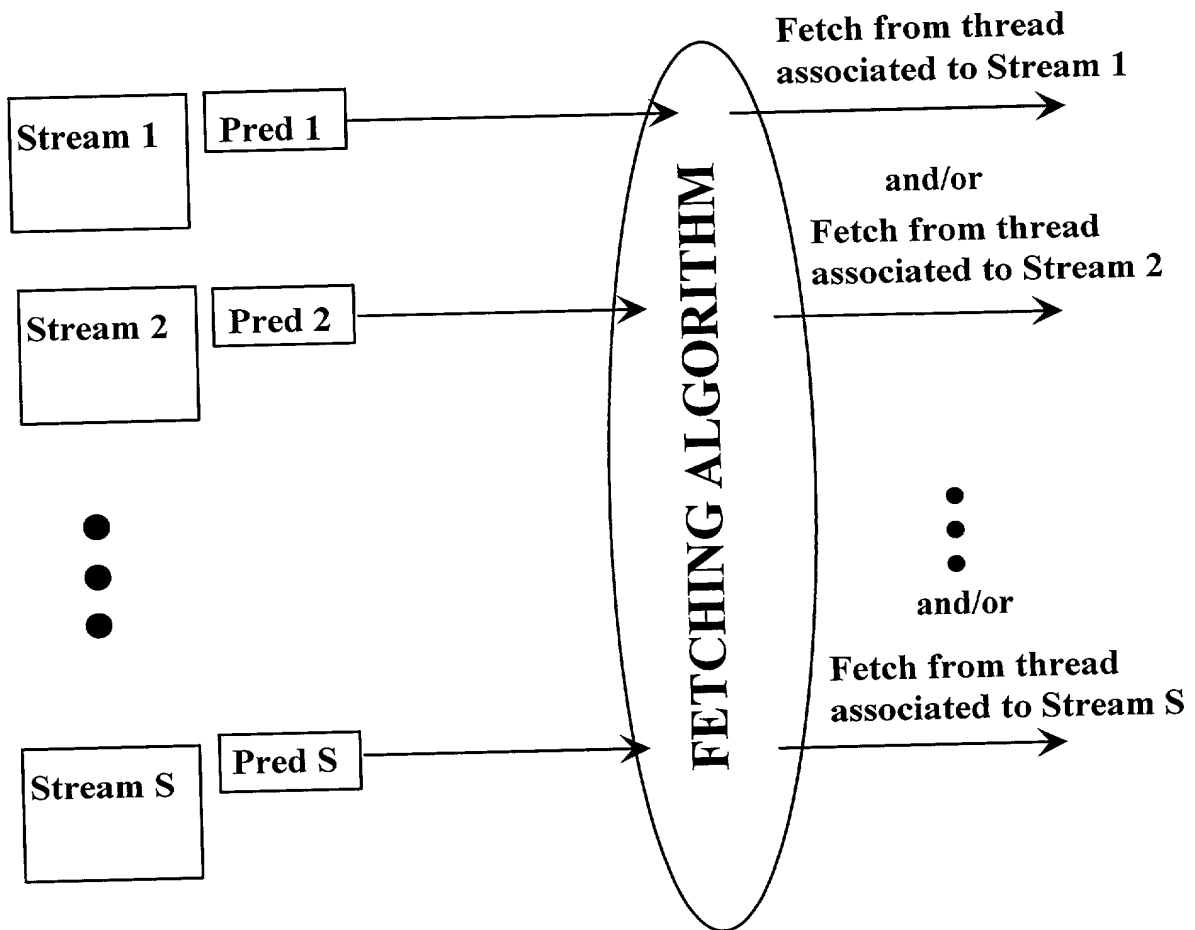


Fig. 2




Fig. 3

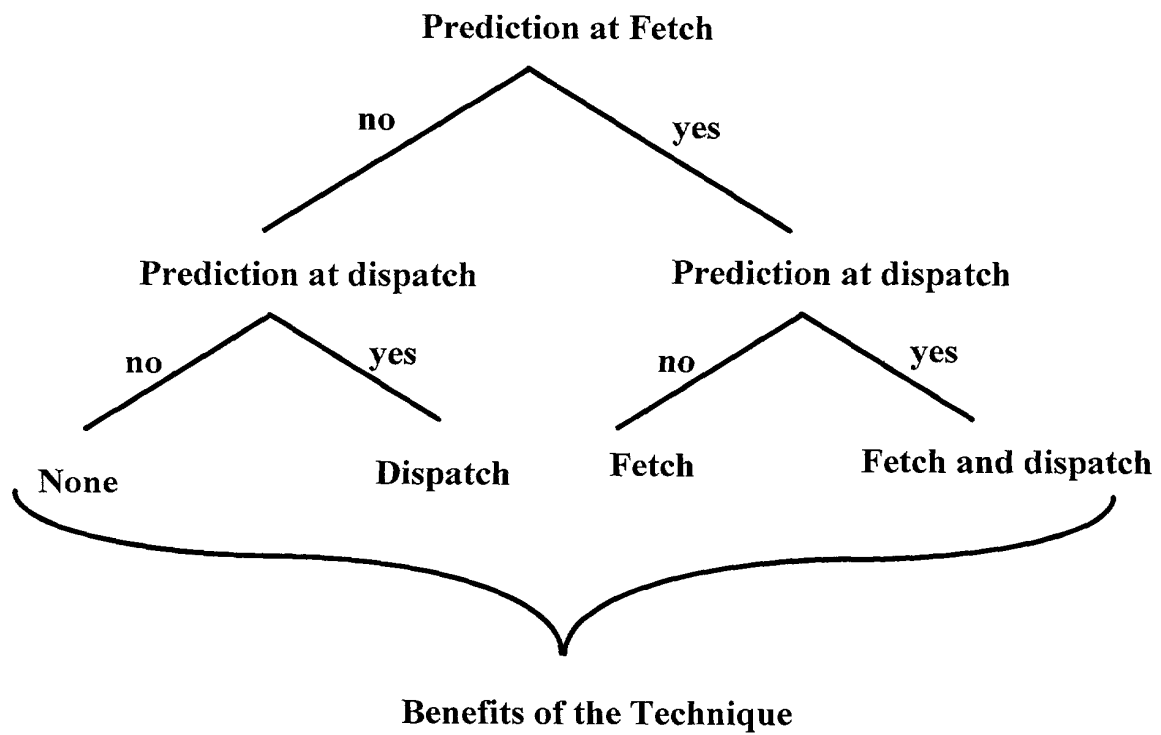


Fig. 4

PC

0 ----

1 ----

2 ----

3 ----

Br

If (x<2) go to 9, else contin.

4 ----

5 ----

6 ----

7 ----

8 ----

9 ----

Fig. 5

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

ATTORNEY DOCKET NO. **P3809**

As a below named inventor, I hereby declare that: My residence, post office address and citizenship are as stated below next to my name. I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: **Methods and Apparatus for Improving Fetching and Dispatch of Instructions in Multithreaded Processors**

the specification of which (check one) ☒ is attached hereto.

☐ was filed on:

☐ Application Serial No.

☐ and was amended on

(If applicable)

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment referred to above. I acknowledge the duty to disclose information which is material to the examination of this application in accordance with Title 37, Code of Federal Regulations, s 1.56 (a). In the case that the present application is a continuation-in-part application, I further acknowledge the duty to disclose material information as defined in 37 CFR s 1.56(a) which became available between the filing date of the prior application and the filing date of the present application. I hereby claim foreign priority benefits under Title 35, United States Code s119 of any foreign applications for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed:

Prior Foreign Application(s)

(Number)

(Country)

(Day/Month/Year Filed)

(Number)

(Country)

(Day/Month/Year Filed)

I hereby claim the benefit under Title 35, United States Code, s120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, s112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, s156(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application.

(Application Serial No.): 09/595,776 (Filing Date): 06/16/2000 (Status): pending

(Application Serial No.): 09/216,017 (Filing Date): 12/16/1998 (Status): pending

(Application Serial No.): 09/240,012 (Filing Date): 01/27/1999 (Status): pending

(Application Serial No.): 09/273,810 (Filing Date): 03/22/1999 (Status): pending

(Application Serial No.): 09/312,302 (Filing Date): 05/14/1999 (Status): pending

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.
(List name and registration number)

Name: Donald R. Boys

Reg. No. 35,074

SEND CORRESPONDENCE TO:

Donald R. Boys

P.O. Box 187

Aromas, CA 95004

DIRECT TELEPHONE CALLS TO:

Donald R. Boys (831) 726-1457

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full name of sole or first inventor: Enric Masoll

1st inventor's signature: [Signature] Dated: 6/15/00
Residence: 7210 Via Romero, San Jose, CA, 95139 Citizenship: Spain
Post Office Address: Same

Full name of 2nd joint inventor, if any: Mario Nemirovsky

2nd inventor's signature: [Signature] Dated: 6/15/00
Residence: Saratoga, CA Citizenship: USA
Post Office Address: 750 University Ave., Los Gatos, CA, 95032

Full name of 3rd joint inventor, if any: _____

3rd inventor's signature: _____ Dated: _____
Residence: _____ Citizenship: _____
Post Office Address: _____

Full name of 4th joint inventor, if any: _____

4th inventor's signature: _____ Dated: _____
Residence: _____ Citizenship: _____
Post Office Address: _____

Full name of 5th joint inventor, if any: _____

5th inventor's signature: _____ Dated: _____
Residence: _____ Citizenship: _____
Post Office Address: _____

Full name of 6th joint inventor, if any: _____

6th inventor's signature: _____ Dated: _____
Residence: _____ Citizenship: _____
Post Office Address: _____

Full name of 7th joint inventor, if any: _____

7th inventor's signature: _____ Dated: _____
Residence: _____ Citizenship: _____
Post Office Address: _____

Full name of 8th joint inventor, if any: _____

8th inventor's signature: _____ Dated: _____
Residence: _____ Citizenship: _____
Post Office Address: _____